

Democracy at Risk: Manipulating Search Rankings Can Shift Voting Preferences Substantially Without Voter Awareness

Robert Epstein (re@aibr.org) & Ronald E. Robertson

American Institute for Behavioral Research and Technology

Abstract

In a double-blind, controlled experiment, web pages and search engine results from an actual election were presented to three groups of eligible voters. In two of the groups, rankings favored one candidate or the other. Preferences shifted dramatically toward favored candidates, with 75% of subjects showing no awareness of the manipulation. In a second experiment, voter preferences again shifted in the predicted direction, and the proportion of people who were unaware of the manipulation was increased by slightly altering the rankings to mask the favoritism. In a third experiment, a more aggressive mask was used to hide the manipulation, and no subjects appeared to be aware of it, even though voter preferences still shifted in the predicted directions. We conclude (1) that the outcomes of real elections—especially tight races—could conceivably be determined by the strategic manipulation of search engine rankings and (2) that the manipulation could be accomplished without people being aware of it.

Detailed Summary

Internet search rankings have a significant impact on consumer choices, mainly because most users click only on highly ranked results (Agichtein, Brill, Dumais, & Ragno, 2006; Granka, Joachim, & Gay, 2004; Guan & Cutrell, 2007; Joachims et al., 2007; Pan et al. 2007). This is why North American companies now spend more than 20 billion dollars annually to place results at the top of rankings (Econsultancy, 2012; Learmonth, 2010). We conducted an experiment to determine whether the deliberate manipulation of search rankings could also influence voter preferences.

In the first of three experiments, we employed a double-blind control-group design with random assignment. 102 subjects were recruited through newspaper ads in the San Diego, California area and were paid a nominal fee for their participation. They were pre-screened in an attempt to match the diverse demographic characteristics of the U.S. voting population. Subjects ranged in age from 18 to 81 with a mean age of 36.6. They reported conducting an average of 8.9 searches per day using search engines, and 67.6% reported having conducted searches to learn about political candidates.

All subjects were shown actual search results and web pages from the 2010 election to determine the prime minister of Australia. This election was used to minimize possible bias by U.S. study participants, and our subjects indeed reported having little or no familiarity with the candidates, Julia Gillard and Tony Abbott (an average familiarity level of 1.3 on a scale of 10).

Subjects were randomly assigned to one of three groups: (1) rankings favoring Gillard, (2) rankings favoring Abbott, or (3) rankings favoring neither (Fig. 1a-c). Neither the subjects nor the research assistants who supervised them knew either the hypothesis of the experiment or the groups to which subjects were assigned. Initially, subjects read brief biographies of the candidates and rated them. Then they spent up to 15 minutes gathering more information about the candidates using a mock search engine showing five pages of search results with six results per page, after which they rated the candidates again. The same web pages and search results were used for all subjects.

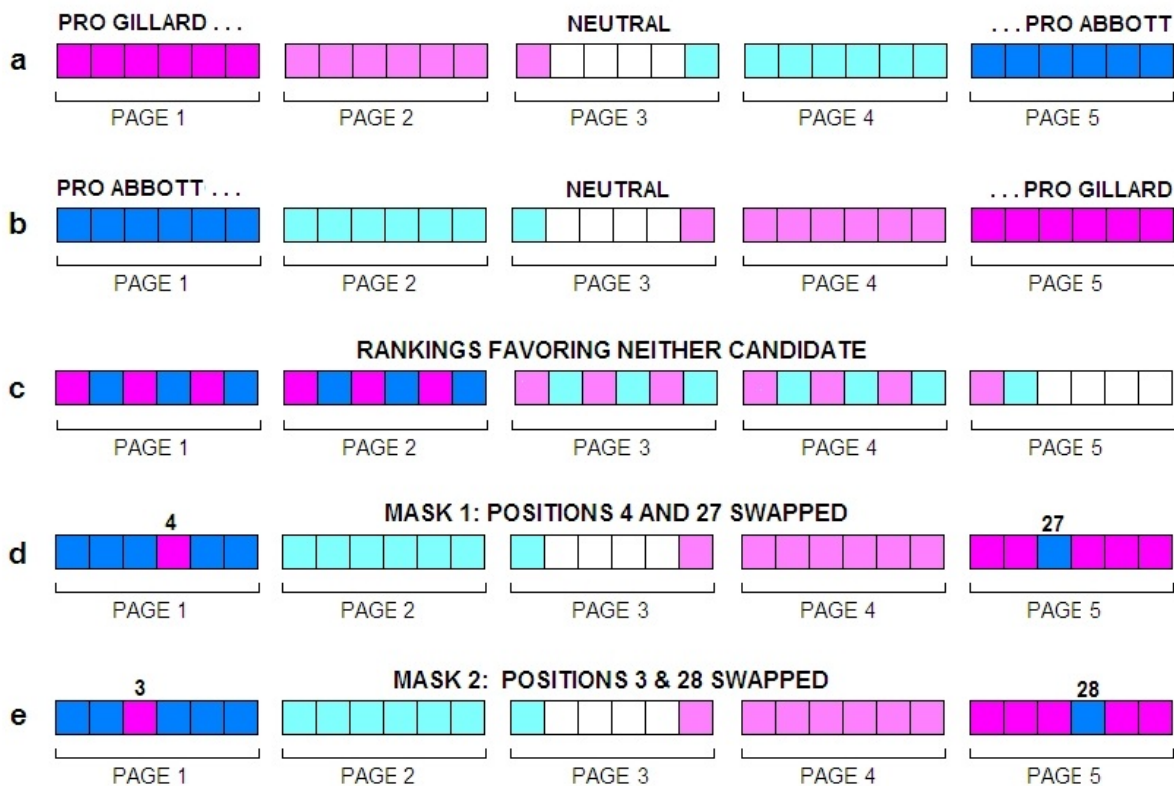


Figure 1 | Search rankings for the three experiments. a, For subjects in Group 1 of Experiment 1, 30 search results that linked to 30 corresponding web pages were ranked in an order that favored candidate Julia Gillard. b, For subjects in Group 2 of Experiment 1, the search results were displayed in precisely the opposite order so that they favored the opposing candidate, Tony Abbott. c, For subjects in Group 3 of Experiment 1 (the control group), the ranking favored neither candidate. d, For subjects in Groups 1 and 2 of Experiment 2, the rankings bias was masked slightly by swapping results that had originally appeared in positions 4 and 27. Thus, on the first page of search results, five of the six results – all but the one in the 4th position – favored one candidate. e, For subjects in Groups 1 and 2 of Experiment 3, a more aggressive mask was employed by swapping results that had originally appeared in positions 3 and 28.

Prior to the web research, no significant differences existed between group ratings. Following the web research, statistically significant differences emerged with respect to how much subjects liked and trusted the candidates and with respect to their overall impressions of the candidates. In addition, voting preferences diverged by 3.71 points on an 11-point scale ($p < 0.01$), and 64.7% and 67.6% of subjects in the bias groups, respectively, said they would vote for the favored candidate, compared to an even 50/50 split in the control group. Notably, 75% of subjects showed no awareness of the manipulation.

In a second experiment, the same procedure was used with 102 additional subjects who were recruited and screened as described above. In this experiment, however, we sought to determine whether the proportion of subjects who are unaware of the manipulation could be increased while voter preferences still shifted in the predicted directions. We accomplished this by slightly altering the order in which the search results had been presented in Experiment 1. Specifically, the search result that had appeared in the fourth position on the first page of the search results favoring Abbott was swapped with the corresponding search result favoring Gillard (Fig. 1d). Thus, the bias in the rankings was now slightly masked.

As a result, voter preferences still shifted in the predicted directions, with the margin between the votes increasing from 3.71 in Experiment 1 to 4.44 in Experiment 2 ($p < 0.01$). In addition, the proportion of people who showed no awareness of the manipulation increased from 75% to 85%, suggesting, perhaps, that a stealthier approach to search ranking manipulations might have more impact than one that is somewhat obvious.

In a third experiment, a more aggressive mask was employed, namely, the search result that had appeared in the third position on the first page of the search results favoring Abbott was swapped with the corresponding search result favoring Gillard (Fig. 1e). Thus, on the first page of search results, five out of six of those results favored one candidate, and the third result favored the other. Again, 102 subjects were recruited, and the procedure was repeated as previously described. Once again, voter preferences shifted in the predicted directions (vote margin 2.64, $p < 0.01$), but this time *no* subjects appeared to be aware of the manipulation.

Real elections are often won by small vote margins. 50% of U.S. presidential elections have been won by vote margins under 7.6%, and 25% of U.S. senatorial elections in 2012 were won by vote margins under 6% (Leip, 2012; Rogers & Cage, 2012). Because races are often so close, search ranking manipulations have the potential to determine the outcomes of many elections in countries worldwide, especially by impacting “swing voters”—that is, voters who remain undecided as the election day grows near. The Australian election in question was won by a margin of only 0.24% and perhaps could easily have been “flipped” by such a manipulation.

The results of the second and third experiments also suggest that, with sufficient study, optimal ranking strategies could be developed that would alter voter preferences while making the ranking manipulations undetectable. Even if a statistical analysis showed that rankings consistently favored one candidate over another, those rankings could always be attributed to algorithm-guided dynamics driven by market forces; deliberate manipulation

would be difficult to prove. Restricting rankings manipulations to voters who have been identified as undecided while also donating money to favored candidates would be an especially powerful and efficient way of wielding influence.

Because search rankings are based on the popularity of websites, voter preferences likely impact those rankings. Given that rankings can in turn affect voter preferences, these phenomena might interact synergistically, causing an explosion of support for one candidate at some point even when the effects of the individual phenomena are small. A mathematical model we have developed shows the possible dynamics of such synergy:

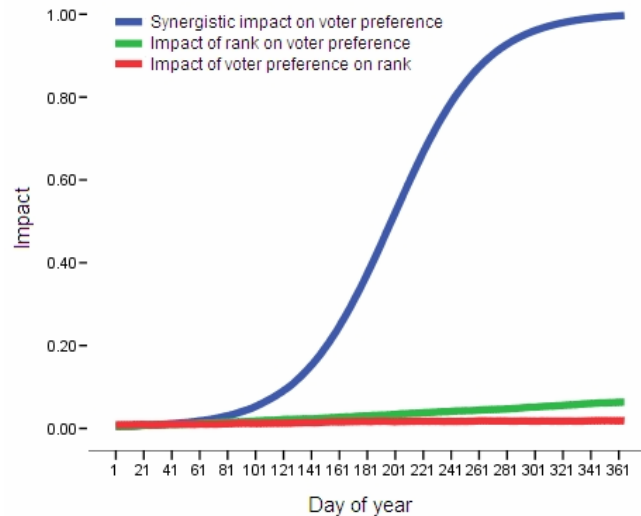


Figure 2 | A possible synergistic relationship between the impact that search engine rankings have on voter preferences and the impact that voter preferences have on search engine rankings. The lower curves (red and green) show slow increases that might occur if each of the processes acted alone over the course of a year (365 iterations of the model). The upper curve (blue) shows the result of a possible synergy between these two processes using the same parameters that generated the two lower curves. The curves are generated by an iterative model employing equations of the general form $V_{n+1} = V_n + r (R_n * (1 - V_n)) + r (O_n * (1 - V_n))$, where V is voter preference for one candidate, R is the impact of voter preferences on search rankings, O is the impact (randomized with each iteration) of other influences on voter preferences, and r is a rate-of-change factor. Because a change in voter preference alters the proportion of votes available, its value in the model cannot exceed 1.0.

Given that search companies are currently unregulated, our results could be viewed as a cause for serious concern, suggesting that such companies could manipulate the outcomes of elections in ways that would be difficult to detect. And whereas donating to campaigns is expensive, altering search results costs nothing.

We acknowledge that voters are subjected to a wide variety of influences during political campaigns: television commercials, political rallies and debates, telephone calls, mailings, and so on. We believe, however, that search ranking manipulations might exert a disproportionately large influence over voters for four reasons:

First, as we have noted, the process by which search rankings affect voter preferences might interact synergistically with the process by which voter preferences affect search rankings, thus greatly magnifying the potential impact of even very subtle search ranking manipulations.

Second, most if not all of the usual sources of influence are explicit and obvious in their allegiance to one candidate or another, whereas ranking manipulations are, in theory, undetectable. Influence that is obvious in its nature is easier to resist than influence that is subtle or, in the extreme case, undetectable (Berger & Fitzsimons, 2008; Brasel & Gips, 2011; Fransen, Fennis, & Pruyn, 2008). What's more, when people are unaware that they are being manipulated, they often feel that the decisions they are making are entirely voluntary (Bargh, Lee-Chai, Barndollar, Gollwitzer, & Trötschel, 2001; Pronin & Kugler, 2007). We are not talking here about "subliminal" influences—that is, influence by stimuli that occur at energy thresholds impossible to perceive (Légal, Chappé, Coiffard, & Villard-Forest, 2012; Karremans, Strobe, & Claus, 2006; Strahan, Spencer, & Zanna, 2002; Winkielman, Berridge, & Wilbarger, 2005). Search results and web pages are easy enough to perceive; it is the *pattern of rankings* that people can't see.

Third, all of the candidates and factions in an election typically have access to and utilize the same forms of influence. Although the better funded campaigns typically bear down more strongly, in theory the playing field is level, with each player trying to exert more influence than the others. In other words, with virtually every mechanism of influence one can imagine, candidates are *competing* for people's attention and support. But with little or no competition among search engines in today's marketplace, search engine manipulations could be slanted almost entirely toward one candidate or party in election after election with no one the wiser. Dominance by one company in the search engine business, combined with the invisibility of the manipulations, could, over time, subvert the mechanisms that maintain open and free elections.

Finally, with the attention of voters gradually shifting toward the Internet and away from traditional sources of information such as newspapers, magazines, and even television (Kohut & Rainie, 2000; Purcell, Brenner, & Rainie, 2012; Smith, 2011), the potential influence of search engine rankings on voter preferences will inevitably grow over time, as will the influence of people who have the power to control such rankings.

We conjecture, therefore, that unregulated search rankings could pose a significant threat to a democratic system of government.

About the authors: Robert Epstein is Senior Research Psychologist at the American Institute for Behavioral Research and Technology (AIBRT) in Vista, California, USA, a nonprofit, nonpartisan research organization. A Ph.D. of Harvard University, Dr. Epstein has published 15 books and more than 200 articles on artificial intelligence and other topics and is the former editor-in-chief of *Psychology Today* magazine. Ronald E. Robertson is a graduate of the University of California San Diego and Assistant Director of AIBRT. Correspondence should be directed to the first author at re@aibr.org.